



Face Recognition in the Era of the Cloud and Social Media: Is it Time to Hit the Panic Button?

by

Joseph J. Atick, Ph.D.

Vice Chairman of IBIA

Face recognition, or the ability of computers to automatically pick up faces in photographs and identify who they belong to, is not new. You saw it in *James Bond* films nearly two decades before the technology was actually invented in academic research centers in the early nineties; and over the last ten years, the technology, propelled by legitimate security needs in the aftermath of 9/11, has evolved dramatically from those fictional depictions to become very real. Today, along with technologies that measure patterns of fingerprints and the iris of the eye, face recognition is a cornerstone in the ensemble of modern *biometrics* which aim to establish individual identity based on the uniqueness of measurable characteristics of the human body.

The applications of these technologies abound in our security-conscious modern society, which is increasingly relying on identity as the foundation for enabling the actions of the honest majority while protecting itself against those who seek to do it harm. Programs that use biometrics – to uniquely identify individuals in order to combat identity fraud, expedite border crossings or dispense social services – are now very common and their scope continues to grow. India, for example, is in the process of enrolling the biometrics of more than 1.2 billion people as part of its social inclusion efforts to issue a unique number (UID) to each individual, which will serve as their unique identity as they interact with their government or their commercial service providers. Biometrics not only protect the integrity of the program by ensuring that no individual enrolls more than once, but they also provide the trust necessary for verifying a claimed identity at a point of service or a transaction thereafter.

Within the family of biometrics currently deployed, face recognition occupies a special position. It is a biometric that can be surreptitiously performed from a distance, without subject cooperation and works from ordinary photographs without the need for special enrollment. These factors, absent responsible-use covenants, raise troubling privacy issues. The industry, recognizing these potential concerns, has sought to self-regulate the application of the technology since 1998, when the leading association representing the industry – the *International Biometrics & Identification Association* (IBIA) – was founded. The IBIA formulated best-practice guidelines which it encouraged its membership and their government and commercial customers to adhere to, providing for the use of the technology in commercial, law enforcement and national security applications in a manner that ensured the continued protection of privacy. As a consequence, the technology evolved in an orderly fashion and its

legitimate applications broadened into a spectrum which today includes protection against identity fraud in programs for issuance of drivers' licenses, passports and travel documents; and, in the hands of investigators, the technology has been a powerful tool credited with solving large numbers of criminal and national security cases and saving innocent lives. In the shadow of its contributions, it is hard for anyone to deny that face recognition has played a positive role in our modern and global society.

So where is the cause for alarm...now?

Simply stated, an unprecedented convergence of several technological developments – a perfect storm – is creating an environment where new kinds of face recognition applications that threaten privacy on a very large scale could emerge over the next decade. Furthermore we see that the control of these applications can no longer be solely in the hands of the industry that created the technology, but will require the active cooperation of social media providers and the IT industry to ensure the continued protection of our reasonable expectations of privacy, without crippling legitimate use of this powerful technology.

To appreciate the perfect storm gathering over the horizon of face recognition, it helps to recall how the technology works. At a schematic level, a face recognition system is essentially made up of three basic elements which mirror the human cognitive process: the “eye,” “brain” and “memory.” The “eye” can be any type of digital camera, video or still, or even a source of digital image files. These images are fed into a computing device running specialized algorithms (the “brain”). The algorithms look for the special patterns which can be discerned as faces. Once a face is detected, its image is extracted from the background and converted into unique mathematical characteristics which capture the individual identity. This mathematical code, often referred to as the *faceprint*, can be compared to a database of faceprints associated with known individuals (the “memory”). The similarity of two faceprints is related to the degree of confidence the algorithm has in the resemblance between the associated faces. Unlike photographs – which are susceptible to changes in lighting, pose, facial expressions, hairdos, etc., – faceprints are reasonably invariant to these factors and hence capture the underlying identity.

Of course, key to the whole identification process is the availability of a repository of known faceprints. This is a database containing images of identified faces and their associated faceprints (“face-memory”), which are used as a reference to identify an unknown face appearing in new images. A face recognition system is only as good as its face-memory and is “face-blind” without one.

Up until now, industry self-regulation controlled the responsible use of face recognition by focusing on strictly controlling the face-memory (i.e., controlling what went into the identification database). For example, IBIA advocated that identification databases built by the police and other government agencies for screening or surveillance applications should only

contain wanted individuals and should be subject to audit to ensure that innocent people were not inadvertently added without documented “probable cause.” So as far as the honest majority was concerned, the systems were face-blind and could not recognize them; consequently their continued privacy was assured.

Today, the implementation of this protection mechanism has become much more difficult, primarily because of the ease with which identification databases can be built from publicly available information in the cloud. Images containing faces and their associated identities (identity-tagged photographs) are proliferating on networking sites, corporate sites and on other generally open web pages. With little effort, an off-the-shelf *web-crawler* program can methodically browse the web and in a short period of time build large databases of identified faces. In fact, more efficient versions of these publicly available software programs are employed by the major search engines which, as a result of their ongoing harvesting over the years, have now accumulated face image databases that in their size dwarf the earth’s population.

Consumers have not helped their privacy cause as they continue to contribute to the swelling size of identification databases through their enthusiastic participation in social media and image sharing sites without seriously understanding the long term privacy implications. Admittedly, most of these sites are supposed to be protected by access control protocols which limit access only to those with login credentials or their “friends.” But no one has examined the consequences of a security breach or what would happen if a government forced social media or image sharing sites to grant access to their massive stacks of identity tagged images.

The net result is that the assemblage of identification databases today is a much simpler process and the potential size of these databases is much larger than what could have been imagined in the days before we had Google and the cloud—the days when face recognition databases were painstakingly built by adding individuals one at a time. By far, this is the first and the most critical factor that has changed from a decade ago. But other converging factors aggravate things.

Second, in the past, face recognition algorithms could not accurately identify millions of individuals from everyday photographs in reasonable computing time. The algorithms were slow and worked best from frontally posed images and under controlled lighting. But anyone who has been following the progress of the technology knows that these and other restrictions are quickly evaporating. Face recognition algorithms, as independently measured by the *National Institute of Standards and Technology (NIST)*, have improved by at least two orders of magnitude over the last decade and have gained many orders of magnitude in speed. This means that today’s state-of-the-art algorithms are at least 100 times more accurate and could be a million times faster than they were ten years ago and are now within the realm where they can function by utilizing every-day photographs. With continuing technological improvements, the clouds of a perfect storm begin to darken.

Third, when face recognition was invented, digital cameras were uncommon; inputting images into computers was not seamless, often requiring multiple manual steps. Today billions of high quality digital cameras are in the hands of consumers. They are found in iPhones, Blackberries, and other mobile devices, and in the ubiquitous snap-and-shoot digital cameras. The potency of which is magnified by the fact they are often networked, so their images can be seamlessly uploaded into computers, or they come with powerful processors capable in themselves of running computationally intensive applications, such as face recognition.

So what does all of this add up to?

The convergence of the above three factors creates a worrisome environment because it opens the door for potentially achieving the unthinkable: the linking of online and offline identities. Basically we now have a proliferation of cameras operating in the real world (offline), along with the proliferation of identity-tagged images on the web (online) and a new generation of powerful face recognition algorithms capable of linking the two. For example, a person photographed by a mobile phone can be identified without their knowledge through face recognition using identity-tagged images harvested over the web. Add to this the powerful data mining capabilities provided by the ever more sophisticated web search engines and we now have the ability to surreptitiously construct a detailed profile of someone we snap with our iPhone walking down the street.

One can see why the current convergence of technologies could pose a threat to privacy. In fact, it fosters the creation of “forensic voyeurism” – using mobile phones, face enabled search engines and publicly accessible images to pierce individual anonymity. While each fragment of information about us on the web may be benign in itself, the danger lies when applications assemble them and link them to our lives offline.

This type of voyeurism could simply be driven by normal social curiosity, deviant behavior or vigilantism. In fact the last has already happened on a small scale in the aftermath of the most recent riots in the United Kingdom, where individuals, dubbed the Google Vigilantes, banded together to attempt to use face recognition software over the cloud to identify the individuals appearing in riot photos posted by ordinary citizens on the web. While their efforts were rudimentary and quickly abandoned, they point to the fact that such behavior will spontaneously emerge should the technical capability become more widely available.

Another disturbing application could emerge when, for example, a special interest social media site teams up with marketing companies interested in identifying site members for the purpose of marketing relevant goods and services to them next time they step into a local store. In what may appear to be a harmless act, the social media site could grant temporary access to its members’ face images from which faceprints could be extracted and geographically segmented. The latter could be aggregated into databases and used by face surveillance

marketing systems in stores to identify and target local shoppers. These systems could even entice shoppers to be identified by rewarding them with special offers or discounts upon identification. In the meantime invaluable information could be extracted about their habits, shopping and otherwise. This is just one example in a whole class of invasive marketing applications that are now in the realm of possibility.

Some may say I am being an alarmist: The technology is still not accurate enough to pervasively allow for the identification of people. But I believe it is only a question of time. I have seen the technology from the day it was born twenty years ago – when I walked into my laboratory, the computer detected me as I passed in front of a camera, compared me against its memory bank of two or three people and haltingly pronounced my name – to today where a person can be matched against millions of faces stored in computer memory in lighting speed and with an accuracy that is more than a million times greater than the original algorithm. The march of progress will only continue and in our lifetime, it is likely that we will see a technology with power and accuracy that could force certain disturbing changes in social behavior, if we do not heed the calls for responsible use.

So what can we do?

Condemning the technology serves no purpose.

Attempting to interfere with its progress is futile.

Banning it is a desperate act.

History is filled with accounts of failed attempts to put the lid on technologies that have a legitimate place in society. Protectionism will most certainly fail once more in our hyper-networked and global society. Face recognition is a tool that has a legitimate role in enhancing security and in combating crime and terrorism and it would be unfortunate to see its responsible use derailed by knee-jerk reactions and irrational fear.

We need to address the root cause of this threat to privacy by focusing on the ease with which identification databases can be built through automated harvesting of identity-tagged images over the web. Of course, we cannot prevent consumers from posting images of their lives or tagging them. So how do we prevent these publicly accessible image sources from being assembled into comprehensive identity databases? We believe there is a series of technical measures that collectively provides the necessary protection.

The campaign must begin by changing the attitude of data handlers (e.g., hosting sites, social media and image sharing sites) towards identity-tagged images. Such images must be treated as identity assets of the consumer (just like other personally identifiable information or PII) and access to them should not be granted for any purpose without notice and consent of

their owners. In addition image stores should be protected to the same security measures used in protecting sensitive information such as financial and health care records. So a security breach in an image site should be viewed just as bad as a breach in say health care records in a hospital. This may seem to be excessive. But unless we begin from this principle today, in twenty years we may be haunted by the cloud's long term memory which will provide the future with identity-tagged faces from the past.

As for identity-tagged face images freely available over the web, we need to implement measures which combat web crawlers and other software agents that automatically harvest them for consolidation into image databases. For example, host websites could block all crawlers unless they originate from an authorized source. To gain authorization, a search engine company would need to reaffirm their privacy commitment before being allowed to crawl into otherwise open websites and collect identity-tagged face images. These privacy assurances should include a commitment not to deploy automated face recognition queries against the resulting repositories, nor, a priori, allow direct access to them by third parties, including governmental agencies.

Recent acquisitions of face recognition technology companies by search engine giants should give the consumer reason for concern absent an explanation of the strategic rationale for these acquisitions and a formal commitment not to use face recognition on their accumulated images. We welcome recent statements by senior executives from these companies in this regard but they stop short of a reassuring official commitment.

Of course, in all of this we cannot forget the role of the central stake holders who ultimately own the images, and that is the consumers. Fundamentally, the workflow for uploading images needs to be adapted to give consumers the option to opt-in or out within a consistent framework for privacy. For example, when a consumer uploads photos, the hosting site could run face detection algorithms to establish if the images contain faces, estimate their resolution and advise of the consequences of uploading at such resolution. Face resolution is very relevant since, as is well known, face recognition algorithms are challenged when the face resolution is low (typically less than 20 pixels between the eyes). This means a thumbnail face image will escape the face engines, and hence is not a threat to privacy, unlike a high resolution version. So such a site could offer the choice to automatically lower the face resolution or upload it as is. As for the use of social images for targeted marketing the default must be explicitly and clearly set to opt-out and the consumer must be made to understand in simple language the implications of changing the privacy policy to opt-in. There are several other measures that need to be explored in order to retain informed choice in the hands of the consumer. What shape these can take should be the outcome of the type of dialogue we are calling for between all the stake holders around this issue.

Ultimately, a relationship of trust between the consumer, the search engine companies, social media and the biometric industry has to emerge before we can be at ease that privacy as

we know it today can be safely maintained into the future. The industries that created the technologies now converging into a perfect storm are capable of implementing technical measures to provide the protections that can safeguard privacy. Government policies should encourage the adoption of appropriate technical protections to safeguard privacy and, at the same time, should ensure that these policies do not cripple technological progress or undermine the legitimate application of technology in the service of society.

Dr. Joseph J. Atick, is the Vice Chairman and the co-founder of the International Biometrics & Identification Association (IBIA). He is a recognized early pioneer in the industry, having been one of the original inventors of face recognition technology and the co-founder and leader of several companies in the identity management industry, including Visionics, the first company to commercialize face recognition and Identix the first multi-biometric company which merged with a leading credentialing company in 2006 to form L-1 Identity Solutions, where he served as the Chief Strategy Officer up until 2011 when the company was acquired by Morpho; and where he currently works in corporate strategy. Dr. Atick holds a Ph.D. in mathematical physics from Stanford University.